UNITED STATES PATENT APPLICATION FOR:

GLOBAL I/O TIMING ADJUSTMENT USING CALIBRATED DELAY ELEMENTS

Inventors:

Chee How LIM

Keng L. WONG

Songmin KIM

GLOBAL I/O TIMING ADJUSTMENT USING CALIBRATED DELAY ELEMENTS

FIELD

The present invention is directed to timing adjustment using delay elements. More particularly, the present invention is directed to a delay arrangement wherein a global I/O timing adjustment is provided via calibrated delay elements.

BACKGROUND

Fig. 1 illustrates a high speed digital system 100 including a drive circuit 110 receiving a signal SIG and outputting a signal to receive circuit 130 via an existing connection circuit 120 and connectors C, such drive circuit 110 and receive circuit 130 being driven by a common clock signal CLK provided along a clock line 140. The components within the drive circuit 110 and receive circuit 130 (e.g., both which may be implemented via IC chips) are manufactured to have sub-micron dimensions and micron spacings between such components, and accordingly, signal propagation time from one internal IC element to another internal IC element is substantially negligible. As a result, the internal IC circuits operate at extremely high speeds, e.g., chips typically now operate with internal clock speeds in excess of 100 MHz. The present invention arises from the problem that external component spacings outside of the ICs (e.g., spacing between IC chips) are not matching the component spacings within ICs making it difficult if not impossible to manage synchronization with respect to downstream signals.

1

In a system, there may be physical limitations as to how closely spaced a drive circuit 110 and a receive circuit 130 can be placed. More specifically, in highly dense systems having a plurality of interconnected printed circuit boards (PCBs) with several tens/hundreds of IC chips, a tremendous number of interconnection lines, numerous connectors and several hundreds/thousands of supporting components (e.g., resistors, capacitors, inductors, etc.), a drive circuit 110 and a receive circuit 130 may need to be spaced at a substantial distance D (e.g., up to ten to fifteen inches) from one another. Resultant signal propagation along the substantial distance D, and especially through connectors C and any existing circuit 120 may cause a propagating signal not to meet a setup time of the receive circuit 130, i.e., cause a synchronization mismatch between the drive and receive circuits.

More specifically, assuming that the signal SIG is processed and output by drive circuit 110 at a time t = 0 (Fig. 2) coincident with a first clock pulse 242 of a 100 MHz (i.e., megahertz) clock having 10 ns (i.e., nanoseconds) clock periods, and does not arrive at an input of receive circuit 130 until 13 ns later, such signal cannot be input into receive circuit 130 upon occurrence of the second clock pulse 244, i.e., it arrives too late at the receive circuit. As a further problem, it is unlikely that such output signal will remain prevalent (i.e., valid) at an input to receive circuit 130 for another 6-7 nS so as to be available for capturing by receive circuit 130 upon occurrence of a third clock pulse 246. Accordingly, a window of availability of the propagated output signal at the input of receive circuit 130 does not match a predetermined setup time window required by receive circuit 130.

In high-speed I/O designs, the timing specifications allow for very little variation. The timing allocation for each component comes from estimates that are susceptible to errors. These allocations are sometimes altered after the design is completed to remedy violations. As the designs become increasingly complex and the design process becomes shorter, it is important to add features that allow corrections after IC chips are connected whenever necessary. These capabilities permit the design to be tuned in the face of uncertainties due to aggressive process scaling as well as ever changing product specifications.

A first solution skews the on-board clock routing to the transmitter and receiver chips with respect to each other once the systematic timing offset is known. The advantage to this solution is that the routing skews are quite constant across manufacturing conditions, but this requires additional board re-designs that slows the design process. In a second solution, on-chip delay buffers are added or removed from the transmitter or receiver chips in the data path to shift the timings. The advantage to this approach is that it does not require board re-designs, but it consumes a lot of space (i.e., in all I/O pad cells). In addition, since the cost of compensating these buffers would be astronomical, these non-compensated buffers will suffer from process, voltage, and temperature (PVT) variations. The delay buffers can be placed into the common clock path to remedy the penalty area. Again, these non-compensated delay buffers suffer from PVT variations that help one timing component, such as 200 ps setup time margin gain, but costs another timing component dearly, such as 400 ps hold time margin loss.

BRIEF DESCRIPTION OF THE DRAWING(S)

The foregoing and a better understanding of the present invention will become apparent from the following detailed description of example embodiments and the claims when read in connection with the accompanying drawings, all forming a part of the disclosure of this invention. While the foregoing and following written and illustrated disclosure focuses on disclosing example embodiments of the invention, it should be clearly understood that the same is by way of illustration and example only and that the invention is not limited thereto. The spirit and scope of the present invention are limited only by the terms of the appended claims.

The following represents brief descriptions of the drawings, wherein:

Fig. 1 is a block diagram illustration of a high speed digital system for background discussion;

Fig. 2 is a clock signal waveform used for description of the high speed digital system illustrated in Fig. 1;

Fig. 3 is a high-level block diagram of an example embodiment of the present invention, including self-calibrating delay cells;

Fig. 4 is a waveform diagram illustrating the timing shift effect of the example embodiment shown in Fig. 3;

Fig. 5 is a timing adjustment table for the example embodiment of the invention shown in Fig. 3;

Fig. 6A is a block diagram of the internal compensation loop in the example embodiment;

Fig. 6B is a waveform diagram of the digital compensation technique in the example embodiment;

Fig. 7 is a graph illustrating the timing shifts at the PLL inputs of the example embodiment;

Fig. 8 is a graph illustrating the setup and hold time window across pins in the example embodiment;

Fig. 9 is a diagram of an example of the delay cell in the example embodiment shown in Fig. 3;

Fig. 10 is a diagram of an example of the delay buffer in the example embodiment shown in Fig. 3;

Fig. 11 is a diagram of an example of a Digital-to-Analog converter in the example embodiment shown in Fig. 3;

Fig. 12 is a state transition diagram of the lock sequence state machine in the example embodiment;

Fig. 13 is a diagram of an example of a lock detector for the self-calibrating delay cell in the example embodiment;

Fig. 14 is a diagram of a lock indicator deglitching circuit in the example embodiment;

Fig. 15 is a table of an example lock range of the self-calibrating delay cell across PVT variations in the example embodiment;

Fig. 16 is a graph illustrating the PLL jitter across taps of the self-calibrating delay cell in the example embodiment; and

Fig. 17 is a graph of the setup and hold time shift in the example embodiment of the invention.

DETAILED DESCRIPTION

Before beginning a detailed description of the subject invention, mention of the following is in order. When appropriate, like reference numerals and characters may be used to designate identical, corresponding or similar components in differing figure drawings. Further, in the detailed description to follow, example sizes/values/ranges may be given, although the present invention is not limited to the same. Example arbitrary axes (e.g., X-axis, Y-axis and/or Z-axis) may be discussed/illustrated, although practice of embodiments of the present invention is not limited thereto (e.g., differing axes directions may be able to be assigned). Still further, the clock and timing signal figures are not drawn to scale, and instead, exemplary and critical time values are mentioned when appropriate. With regard to description of any timing signals, the terms assertion and negation may be used in an intended generic sense. More particularly, such terms are used to avoid confusion when working with a mixture of "active-low" and "active-high" signals, and to represent the fact that the invention is not limited to the illustrated/described signals, but could be implemented with a total/partial reversal of any of the "active-low" and "active-high" signals by a simple change in logic. More specifically, the terms "assert" or "assertion" indicate that a signal is active independent of whether that level is represented by a high or low voltage, while the terms "negate" or "negation" indicate that a signal is inactive. As a final note, well known

6

power/ground connections to ICs and other components are not shown for simplicity

of illustration and discussion, and so as not to obscure the invention. Further,

arrangements may be shown in block diagram form in order to avoid obscuring the

invention, and also in view of the fact that specifics with respect to implementation of

such block diagram arrangements are highly dependent upon the platform within

which the present invention is to be implemented, i.e., such specifics should be well

within purview of one skilled in the art. Where specific details (e.g., circuits,

flowcharts) are set forth in order to describe example embodiments of the invention,

it should be apparent to one skilled in the art that the invention can be practiced

without, or with variation of, these specific details.

Although example embodiments of the present invention will be described in

an example computer system and environment, practice of the invention is not

limited thereto, i.e., the invention may be able to be practiced with other types of

systems, and in other types of environments (e.g., communications chips).

Turning now to detailed description, the example embodiments of the

invention use a tuning feature after IC chips have been connected in a system

design, such as on a printed circuit board of a personal computer (PC) system. It

allows a set of delay elements, such as self-calibrating delay cells, to globally shift

I/O timings such that transmission and reception timing characteristics can be

changed with respect to an IC chip. This methodology works in a common-clock

design where both the transmitter and receiver chips are synchronized by the same

clock driver. For example, for a given bus topology, if there is a systematic setup

time violation at the receiver, the self-calibrating delay cells can be activated to delay the receiver's clock path with respect to its data path to eliminate this violation.

The example embodiments have a main technical advantage over other approaches discussed in the Background. The delay cells are self-calibrated using a digital compensation technique that reduce PVT variation without incurring large area penalty. This allows post-silicon adjustments to eliminate systematic timing violations without suffering from overall timing window degradation.

Fig. 3 illustrates the invention from a high-level perspective. In this block diagram, calibrated delay elements (labeled as CDE) are strategically located in the reference and feedback clock paths of a phase locked loop (PLL) 300. As a result, the input and output timings can be systematically shifted with respect to the external clock (xclk). Since PLL 300 provides the timing reference for the entire IC chip through 1/N frequency divider 302 (only a single Input Flip-Flop 303 and Output Flip-Flop 304 are shown in Fig. 3 for the sake of convenience), adding delay elements (T_{ext}) 301-1 in the path of the reference clock signal shifts the internal coreclk and bclk signals later with respect to the external xclk signal. This increases setup time margin (as shown in Fig. 4) when the chip is in the receive mode. However, since the bus period is constant, improving setup (T_{su}) margin in the input path also improves minimum clock-to-output ($T_{co,min}$) margin, but at the cost of hold (T_h) and maximum clock-to-output ($T_{co,max}$) margins in the input and output paths, respectively.

The calibrated delay elements 301-1 and 301-2 allow the margins to be shifted without growing the timing window (T_{su} + T_h input window or $T_{co,min}$ + $T_{co,max}$

output window). In present day high-speed I/O design, growing the timing window hurts the overall timing balance, and translates the timing violation from one component to another. There are some guidelines as to the situation when global timing adjustments are appropriate as shown in Fig. 5.

Fig. 6A shows the block diagram of a preferred self-calibrating delay cell utilizing a digital compensation technique that keeps the delay cells from PVT variation (i.e. providing a continuous calibration mechanism). In this technique, the feedback (fbclk) clock pulse from Fig. 3 is used as the reference signal to the compensation circuit. The falling edge of this signal is phase-aligned to the delayed rising edge (x). To align the edges, the output of the phase detector (PD) 601 tells the Up/Down counter 602 whether the x is early or late with respect to the fbclk falling edge. If early, Up/Down counter 602 will increment its binary code, and the digital-to-analog (DAC) converter 603 will produce a higher voltage that increases the delay of the delay buffer 600. If late, Up/Down counter 602 will increment its value and hence causes the delay to reduce. This process repeats continuously to phase-align the edges. The phase alignment is illustrated in Fig. 6B, and the dotted lines in Fig. 6A indicate the delay taps that are available to provide the fine delay granularity for I/O tuning. As an example, the delayed fbclk signal, which feeds to PLL 300, is shown to be delayed by 4 buffers, while the phase alignment occurs with 5 buffers. Since this circuit is simplistic in nature and highly digitized, the area cost is small, and it only occurs in only one location of the die.

9

Example implementations of various parts of the circuit is illustrated in Figs. 9-11. However, the circuit is of course not limited to such an implementation. Fig. 9 illustrates an example of a single delay cell in the stages 600-1 to 600-n of delay buffer 601. The delay cell 900 receives the P/N bias from DAC 603 in Fig. 6. (An example of DAC 603 is shown in Fig. 11.) Fig. 10 illustrates how a plurality of such delay cells 900 (shown with a diagonal arrow to indicate that they are variable to adapt to, e.g., PVT variations. As shown, there are a plurality of selectable taps corresponding to stages 600-1 to 600-n of delay buffer 600.

Because of the loop circuit in the example self-calibrating delay cell, there is the potential for tuning errors. There will be a certain amount of jitter across the taps of delay buffer 600 as shown in Fig. 16. A lock detector and a lock indicator deglitching circuit may be provided to ensure that the loop circuit does not become stuck at an improper value. Fig. 12 shows a state transition diagram of a lock sequence state machine for the loop circuit. Fig. 13 shows an example of a lock detector for delay buffer 600. Fig. 14 shows an example of a lock indicator deglitching circuit. Of course, there are specified operating conditions within which the circuit must operate and Fig. 15 is a table showing bias codes for examples of the lock range across PVT variations.

Fig. 7 shows the resulting timing shifts at the inputs of PLL 300. The reference clock signal at the center of Fig. 7 is shown with the feedback clock signals shifted in 130 ps granularity. Fig. 8 shows that the pin-to-pin setup and hold timing window remains constant with a timing shift of 2 delay cells (260 ps). These

10

two figures show that I/O timing can be shifted without affecting the timing window. Fig. 17 shows the setup and hold time shift in the example embodiments.

Although a variety of different embodiments are described above, they all provide a flexible and cost-effective way to adjust I/O timings to meet product timing specifications after IC chips have been mounted. The placement of the delay cells only in the input paths of PLL 300 shifts the global timing of the chip with little impact on the amount of area available on an IC chip for other components. In particular, the delay cells calibrate themselves to meet specified timing adjustment granularity and range.

Although not shown for the sake of simplicity, the example embodiments may be implemented in a system including the similar components as Fig. 1. Indeed, the disclosed embodiments and other embodiments of the present invention may be practiced in all types of systems, including, but not limited to, computing systems, non-computing systems, communication systems, etc. The IC chip may be any kind of chip with I/O requirements, including, but not limited to, microprocessors, north bridge, south bridge, memory controller hub, I/O controller hub, an application specific integrated circuit (ASIC), a data interface buffer (DIB) acting as both a transmitting and receiving circuit, and a dynamic random access memory (DRAM) or dual in-line memory module (DIMM) (or other memory type) acting as both a receiving and transmitting circuit.

In actual practice, there may be a single PCB or multiple interconnected PCBs or a multi-layer PCB in a system (such as a server) having an extremely complex system having a layout and components which dictate the spacing between

11

transmitting/receiving IC pairs. There may be a PCB component which is a connector that intervenes and prevents the transmitting device and receiving device from being placed any closer together. Further, a signal propagation path between the devices may vary during the design process, taking into consideration intervening components such as a connector.

While a trend in the art has been to attempt to minimize distances between ICs, the present invention takes a non-obvious approach of increasing an effective signal propagation distance between ICs, i.e., adds delay to the clock signal propagation path synchronizes the drive circuit 110 (in a transmitting device) and a receive circuit 130 (in a receiving device) in order to provide phase delayed synchronization such that downstream signal management is improved. Without phase delayed synchronization, downstream signal management may not be possible due to difficult or impossible management of valid data input timing requirements such as setup and hold times.

The example embodiments of the present invention allow longer propagation paths (i.e., PCB conduction line) while still providing signal propagation match (phase delay synchronize) between the transmitting device and the receiving device. That is, they provide a timing adjustment such that a signal arrival and availability of the signal at a receiving circuit input matches valid data timing input requirements of the receiving circuit.

As a result of the example embodiments, there is little need to minimize PCB spacing distances between clock and transmitting/receiving c12 circuit pairs, and

accordingly, design of complex systems becomes easier as there is more freedom to move sending/receiving components apart to greater separation distances. Second, since there is a direct correlation between PCB conduction line length and delay (e.g., 12 inches of PCB conduction line length ≈ 2 ns of phase synchronization delay), design of complex systems becomes easier. Third, since great numbers of further components and their corresponding clock lines and power connection lines are avoided, the system is less complex and it is less likely that multi-layer PCBs will be required.

In concluding, reference in the specification to "one embodiment", "an embodiment", "example embodiment", etc., means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the invention. The appearances of such phrases in various places in the specification are not necessarily all referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with any example embodiment, it is submitted that it is within the purview of one skilled in the art to effect such feature, structure, or characteristic in connection with other ones of the embodiments. Furthermore, for ease of understanding, certain method procedures may have been delineated as separate procedures; however, these separately delineated procedures should not be construed as necessarily order dependent in their performance, i.e., some procedures may be able to be performed in an alternative ordering, simultaneously, etc.

This concludes the description of the example embodiments. Although the present invention has been described with reference to a number of example embodiments thereof, it should be understood that numerous other modifications and embodiments can be devised by those skilled in the art that will fall within the spirit and scope of the principles of this invention. More particularly, reasonable variations and modifications are possible in the component parts and/or arrangements of the subject combination arrangement within the scope of the foregoing disclosure, the drawings and the appended claims without departing from the spirit of the invention. In addition to variations and modifications in the component parts and/or arrangements, alternative uses will also be apparent to those skilled in the art.

What is claimed is:

14